Memfault

Wrangling Penguins: Better Embedded Linux Monitoring and Debugging with Memfault

Thomas Sarlandie, Linux Tech Lead

Thomas Sarlandie

Linux Tech Lead, Memfault

- Passion: building at the intersection of software and hardware
- Previously led software teams at Pebble and Fitbit
- 🦀🦀🦀 Rust-aficionado

pebble fitbit Memfault

Agenda

- Monitoring embedded linux devices
- Monitoring a fleet of embedded devices
- Debugging with logs and coredumps
- Q & A

Poll #1

Which of these tools do you use to monitor and debug your fleet in production?

Check all that apply

A. SSH

B. Grafana

C. Coredumps

D. Logs

Memfault for Device Reliability Engineering

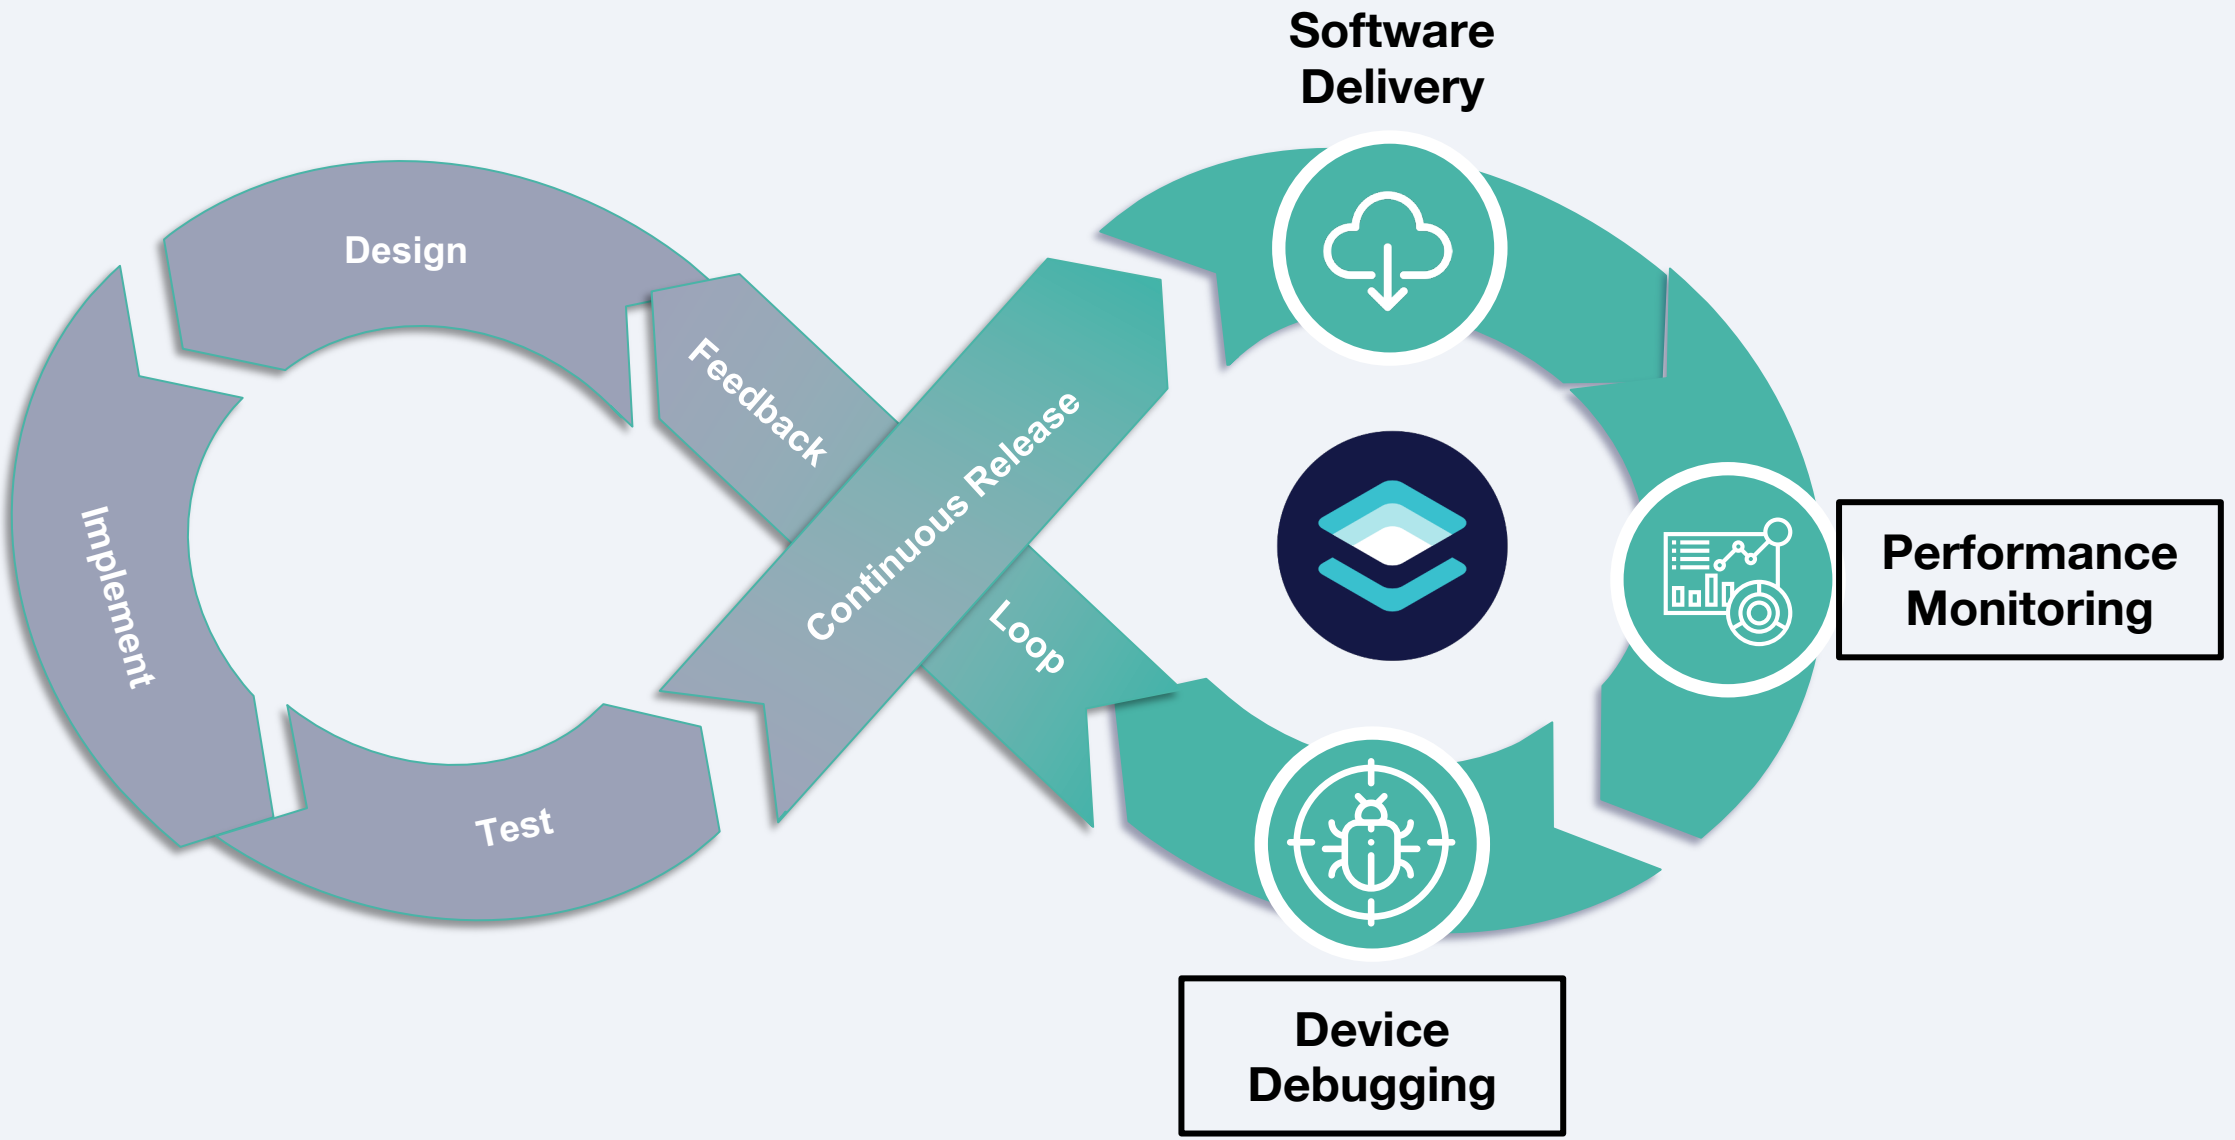

Monitoring embedded linux devices

Monitoring Goals

Validate hypotheses and debug device issues

Get a pulse on the fleet - especially when shipping new hardware or firmware

Detect problems before the customers

Monitoring Challenges

On Device

- Collecting from different sources and languages
- Partial connectivity
- Flash wear and networking costs

Backend

- Scaling pains
- Lack of flexibility
- Visualization tools

Usage

- Drowning in data
- Metrics are meaningless when aggregated
- Signal lost in the data

Metrics

A **metric** is a measurement captured at runtime

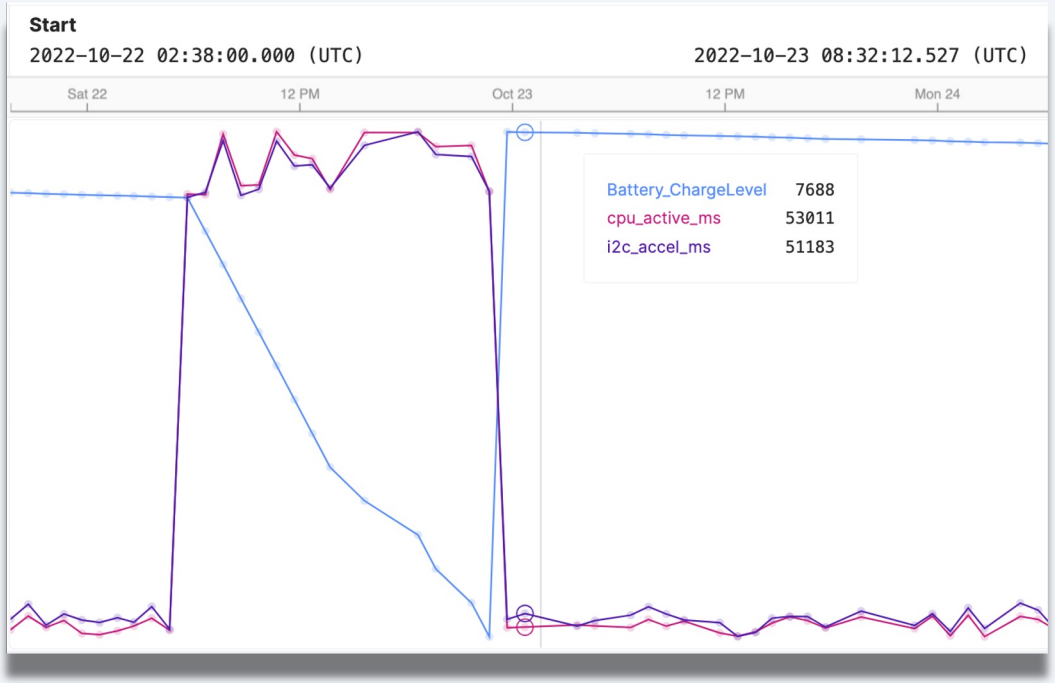

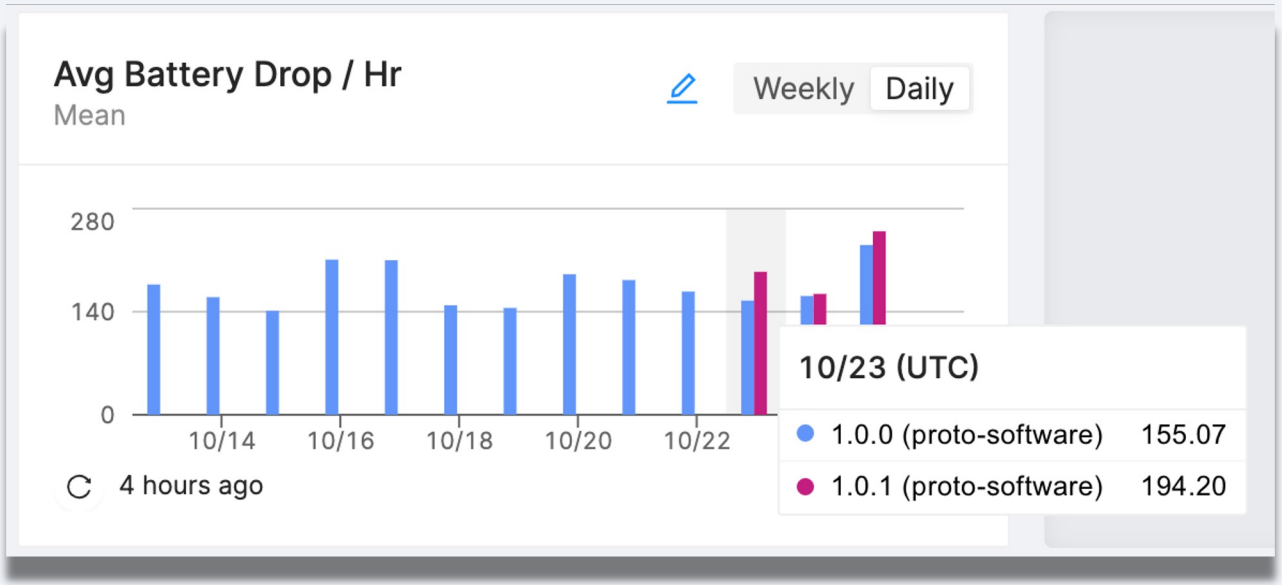

Combing large numbers of metrics and calculating statistics is called an **aggregation**

Collecting metrics on device

- Memfault leverages collectd to capture **system** and **device** metrics
- Customize which system metrics to capture using **collectd plugins**
- Push device metrics using the **collectd/statsd** endpoint

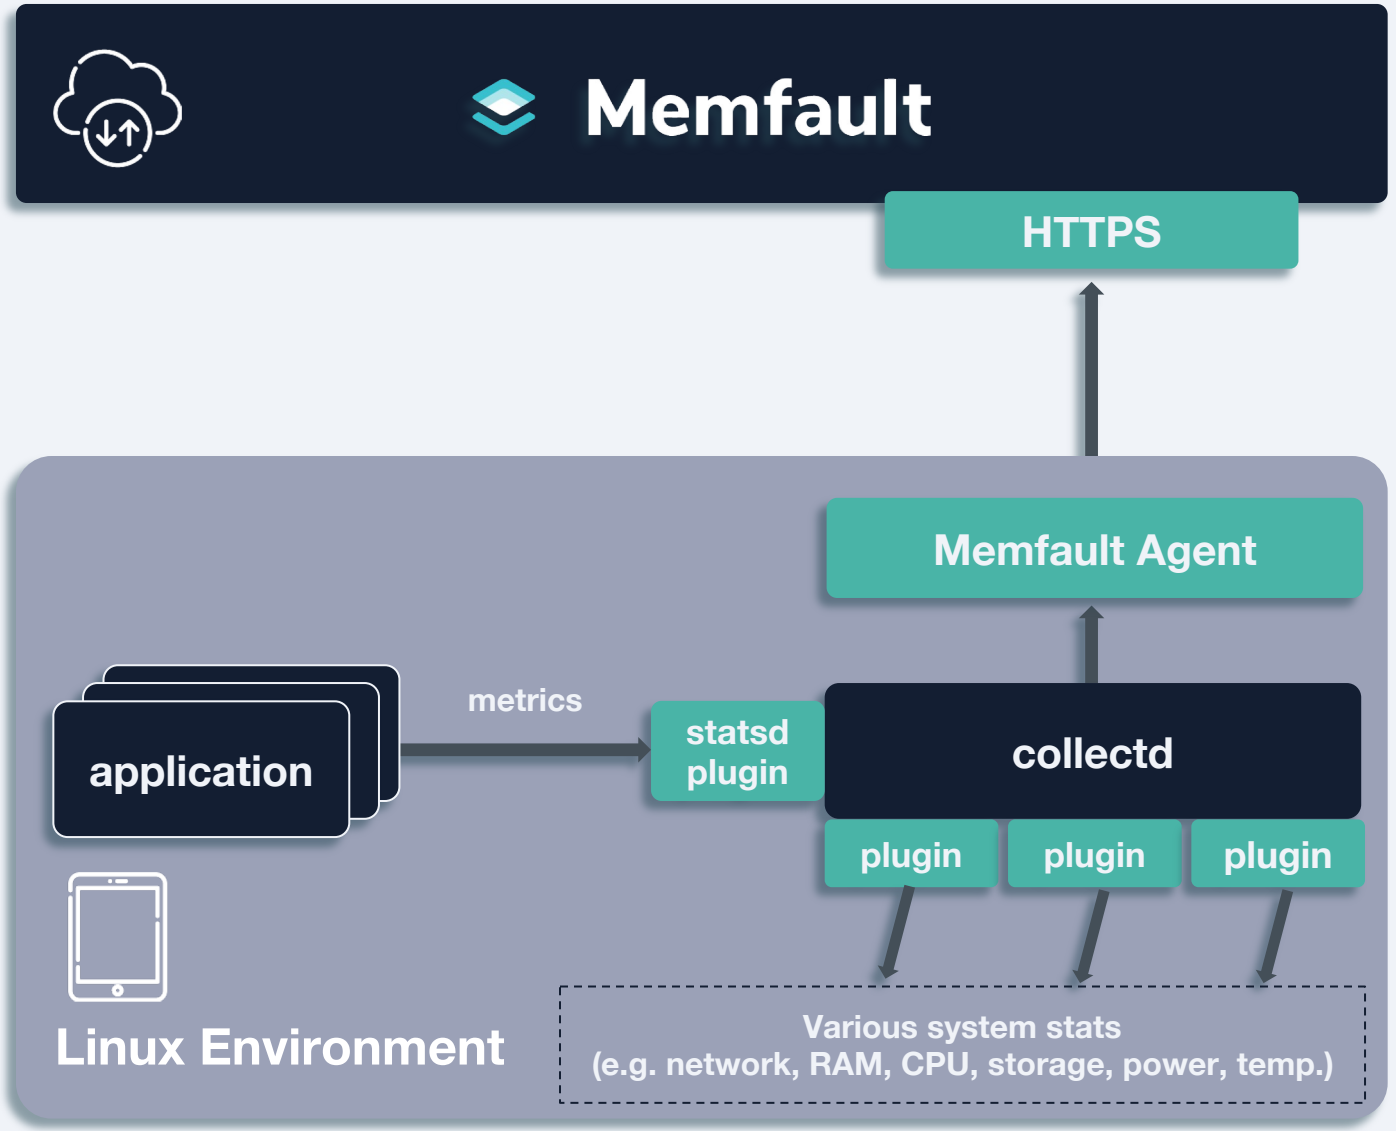

Pushing custom metrics

```c
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <statsd-client.h>
#include <unistd.h>

#define MAX_LINE_LEN 200
#define PKT_LEN 1400

int main(int argc, char *argv[])
{
  statsd_link *link;

  link = statsd_init_with_namespace("localhost", 8125, "mycapp");

  char pkt[PKT_LEN] = {'\0'};
  char tmp[MAX_LINE_LEN] = {'\0'};

  statsd_prepare(link, "mygauge", 42, "g", 1.0, tmp, MAX_LINE_LEN, 1);
  strncat(pkt, tmp, PKT_LEN - 1);
  statsd_send(link, pkt);

  statsd_finalize(link);
}
```

```python
from statsd import StatsClient

statsd = StatsClient(
    host="localhost",
    port=8125,
    prefix="mypythonapp",
)

statsd.gauge("mygauge", 42)
```

Look for a statsd library in your language

Data aggregation that can scale

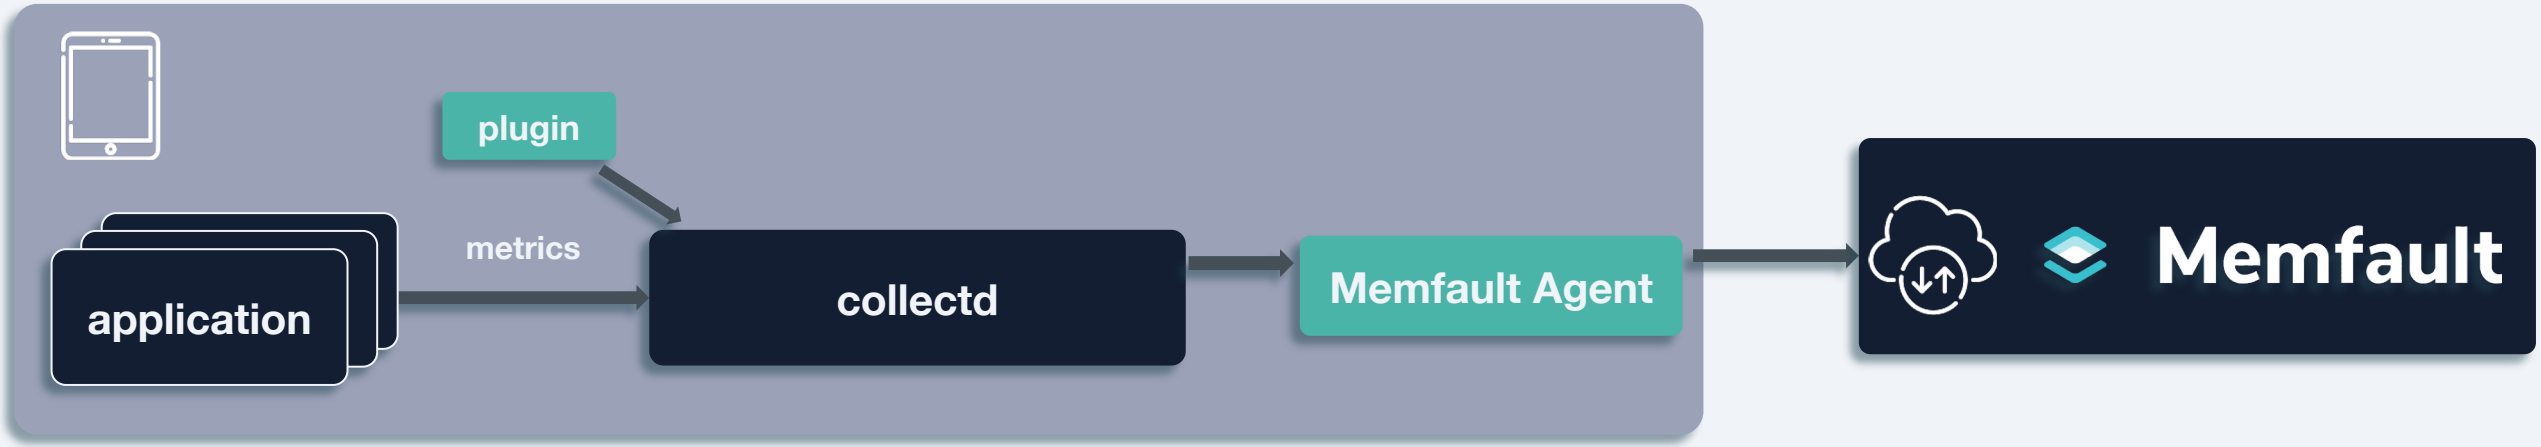

statsd push
Arbitrary frequency

collectd plugins
Typically 10s interval

Collectd Aggregation
Push all metrics every 10s

Memfaultd Aggregation
One heartbeat /hour

Backend Aggregation
Maintain /hour and /day aggregation for all timeseries

Device monitoring

Use metrics to create a device set

- Device sets are dynamic list of devices
- The list will update as new data comes in

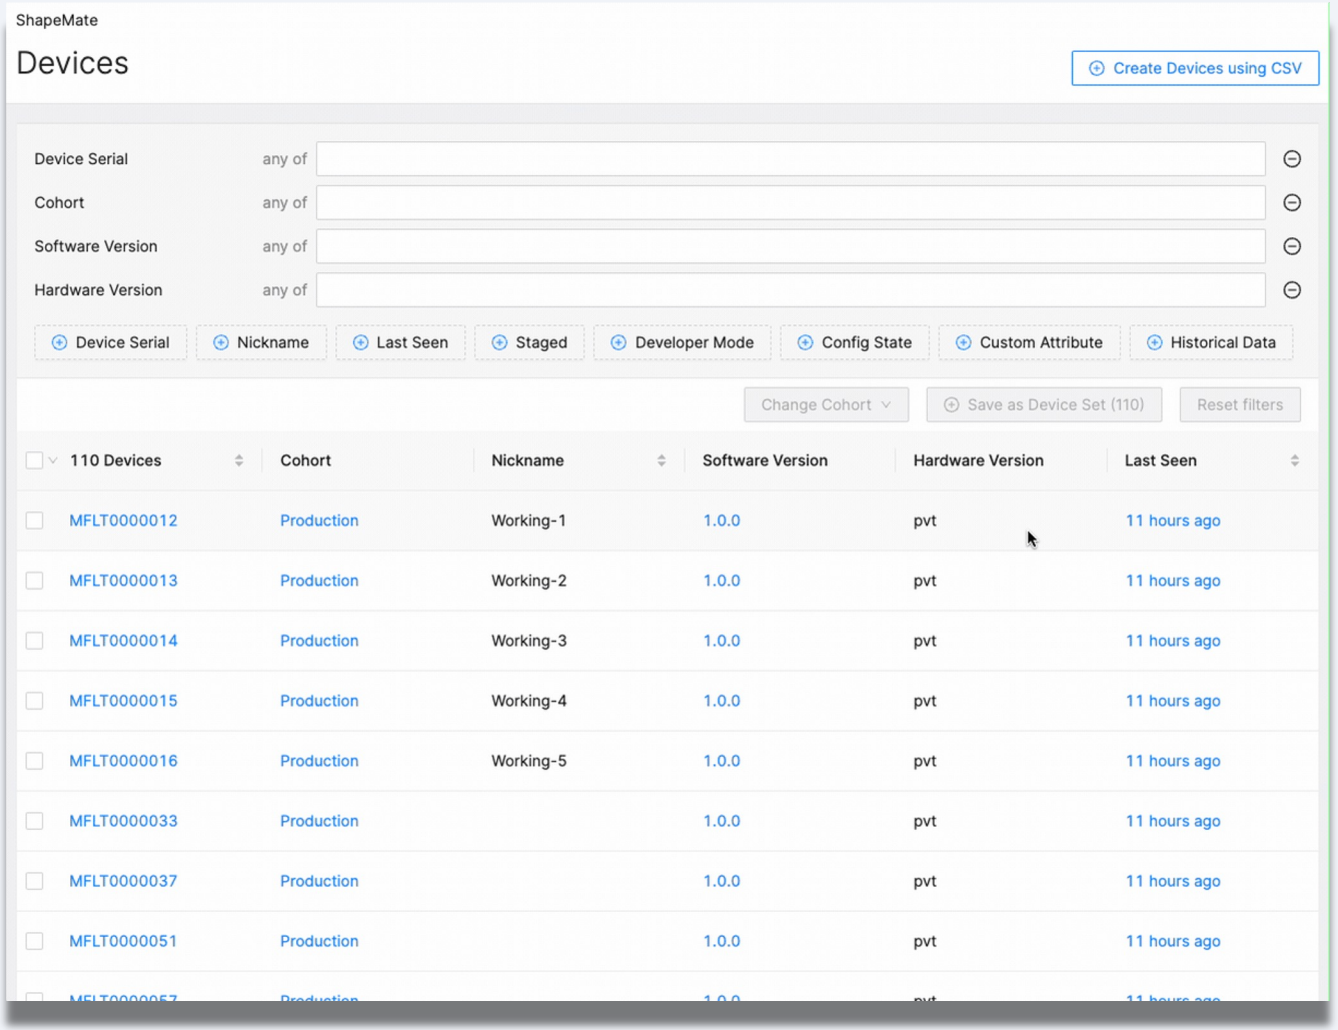

Metrics for alerting

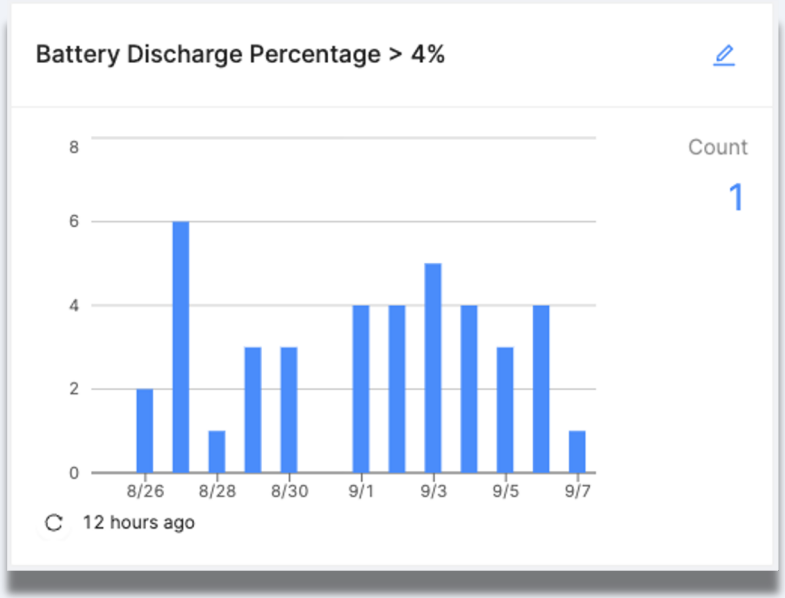

- Device sets make very useful graphs
- Metrics can be used to trigger alerts

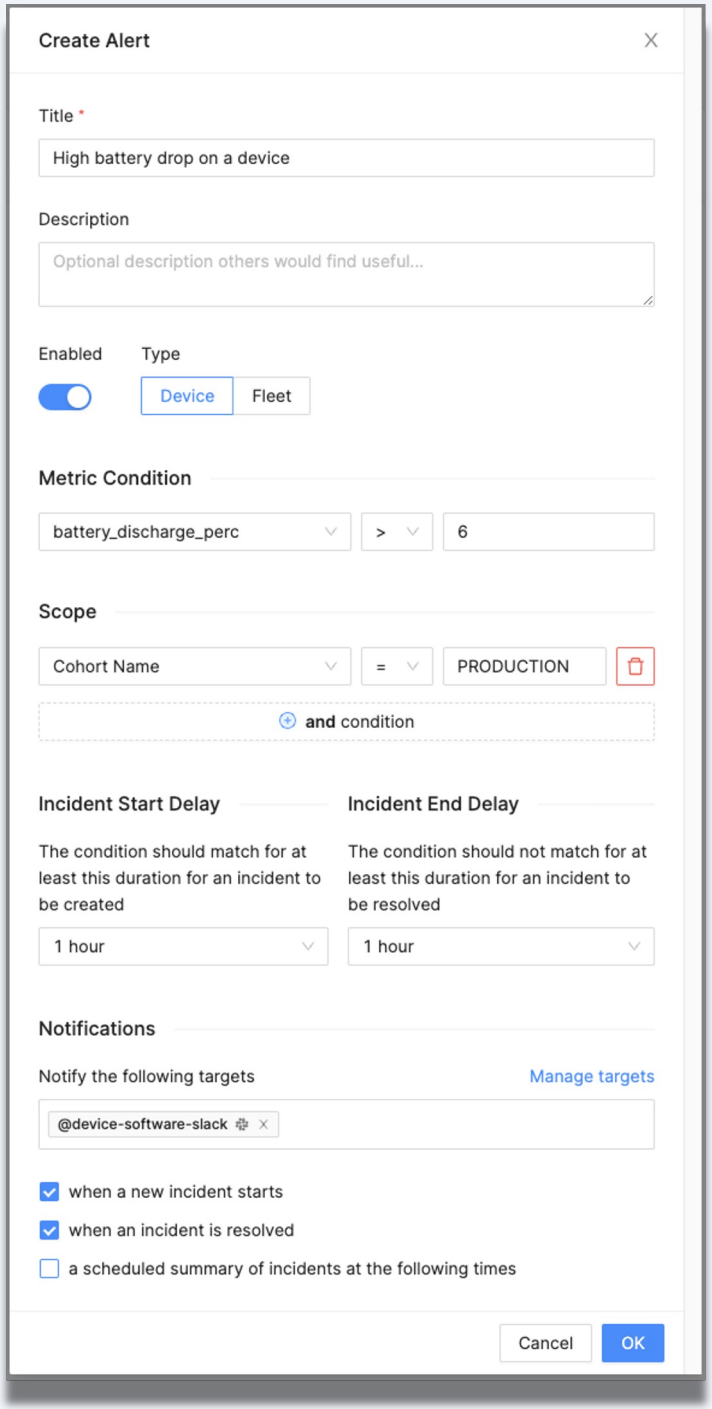

From device monitoring to fleet monitoring

Fleet monitoring

(1000 devices or more)

- Select the right metrics
- Scale the ingestion and aggregation
- Pick the right visualization

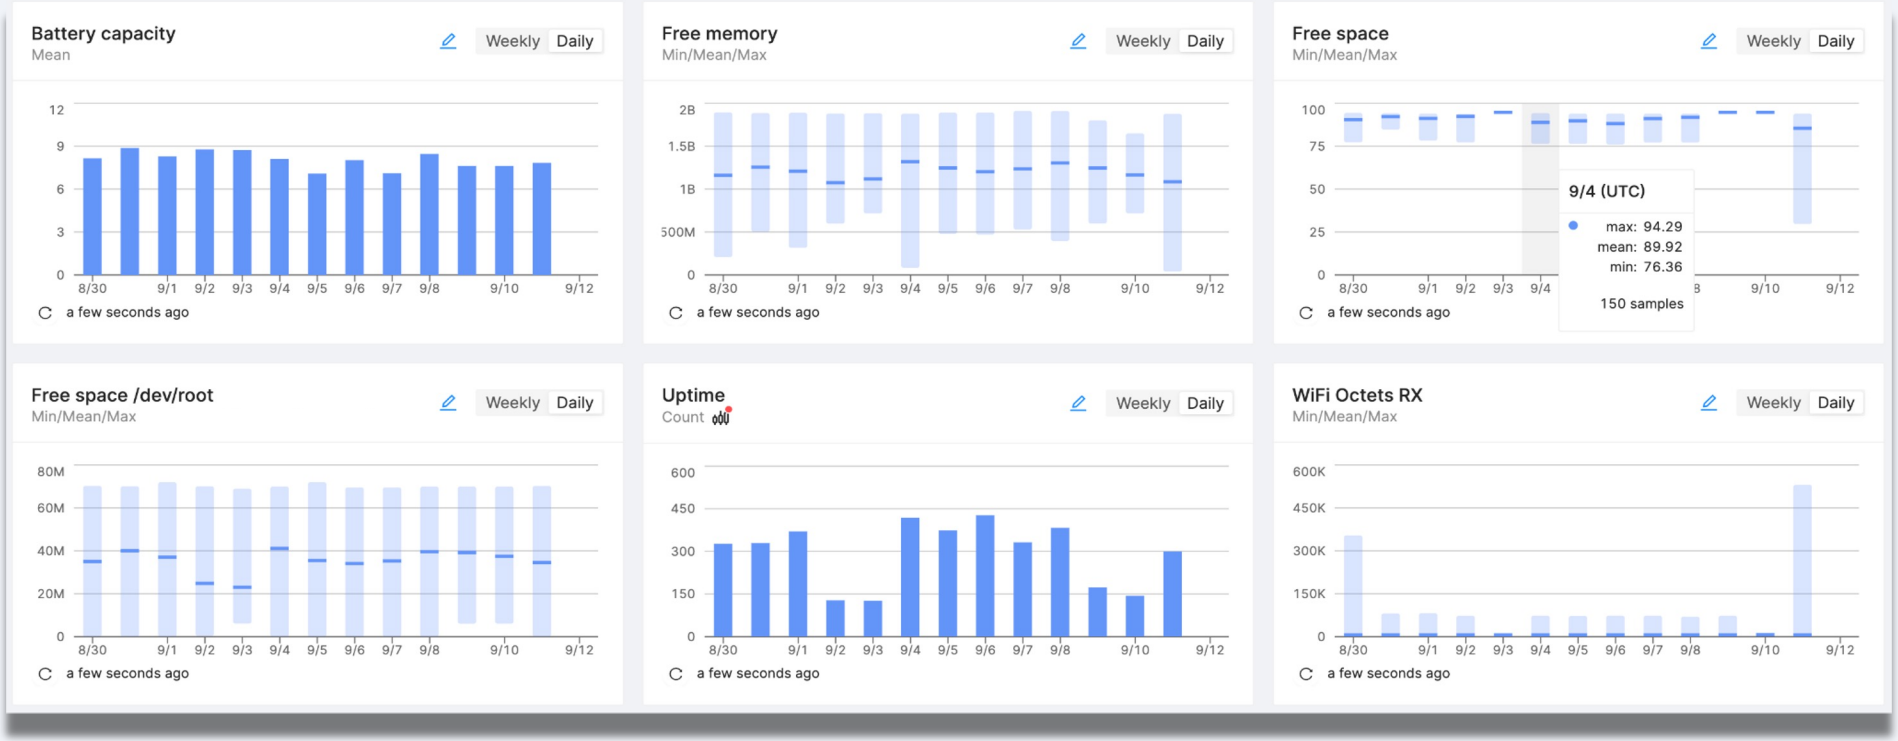

From device to fleet monitoring

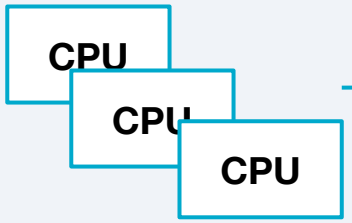

Avg Fleet CPU

❌

Not very useful
unless usage is stable

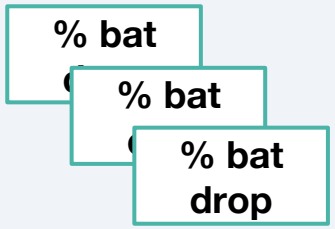

Avg bat %drop per hour

✅

Useful to keep track of power management improvement

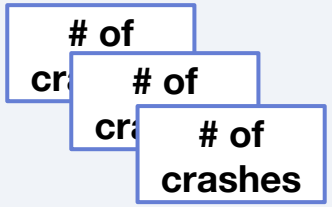

Crash free hours

✅

A useful KPI for device engineering teams

See detailed how to setup crash free hours
https://docs.memfault.com/docs/best-practices/fleet-reliability-metrics-crash-free-hours/

Some useful fleet metrics

Battery

- Discharge per hour (s)
- Screen “on” time (s)

Connectivity

- Cellular modem connected (s)
- Bytes received / sent
- Time spent connecting
- Ping
- RSSI

Flash

- Flash read/write
- Free disk space

Usage

- # of operations / failure
- Time to execute operation

System

- Free memory (Min and Avg)
- Temperature

Visualizing Fleet Data

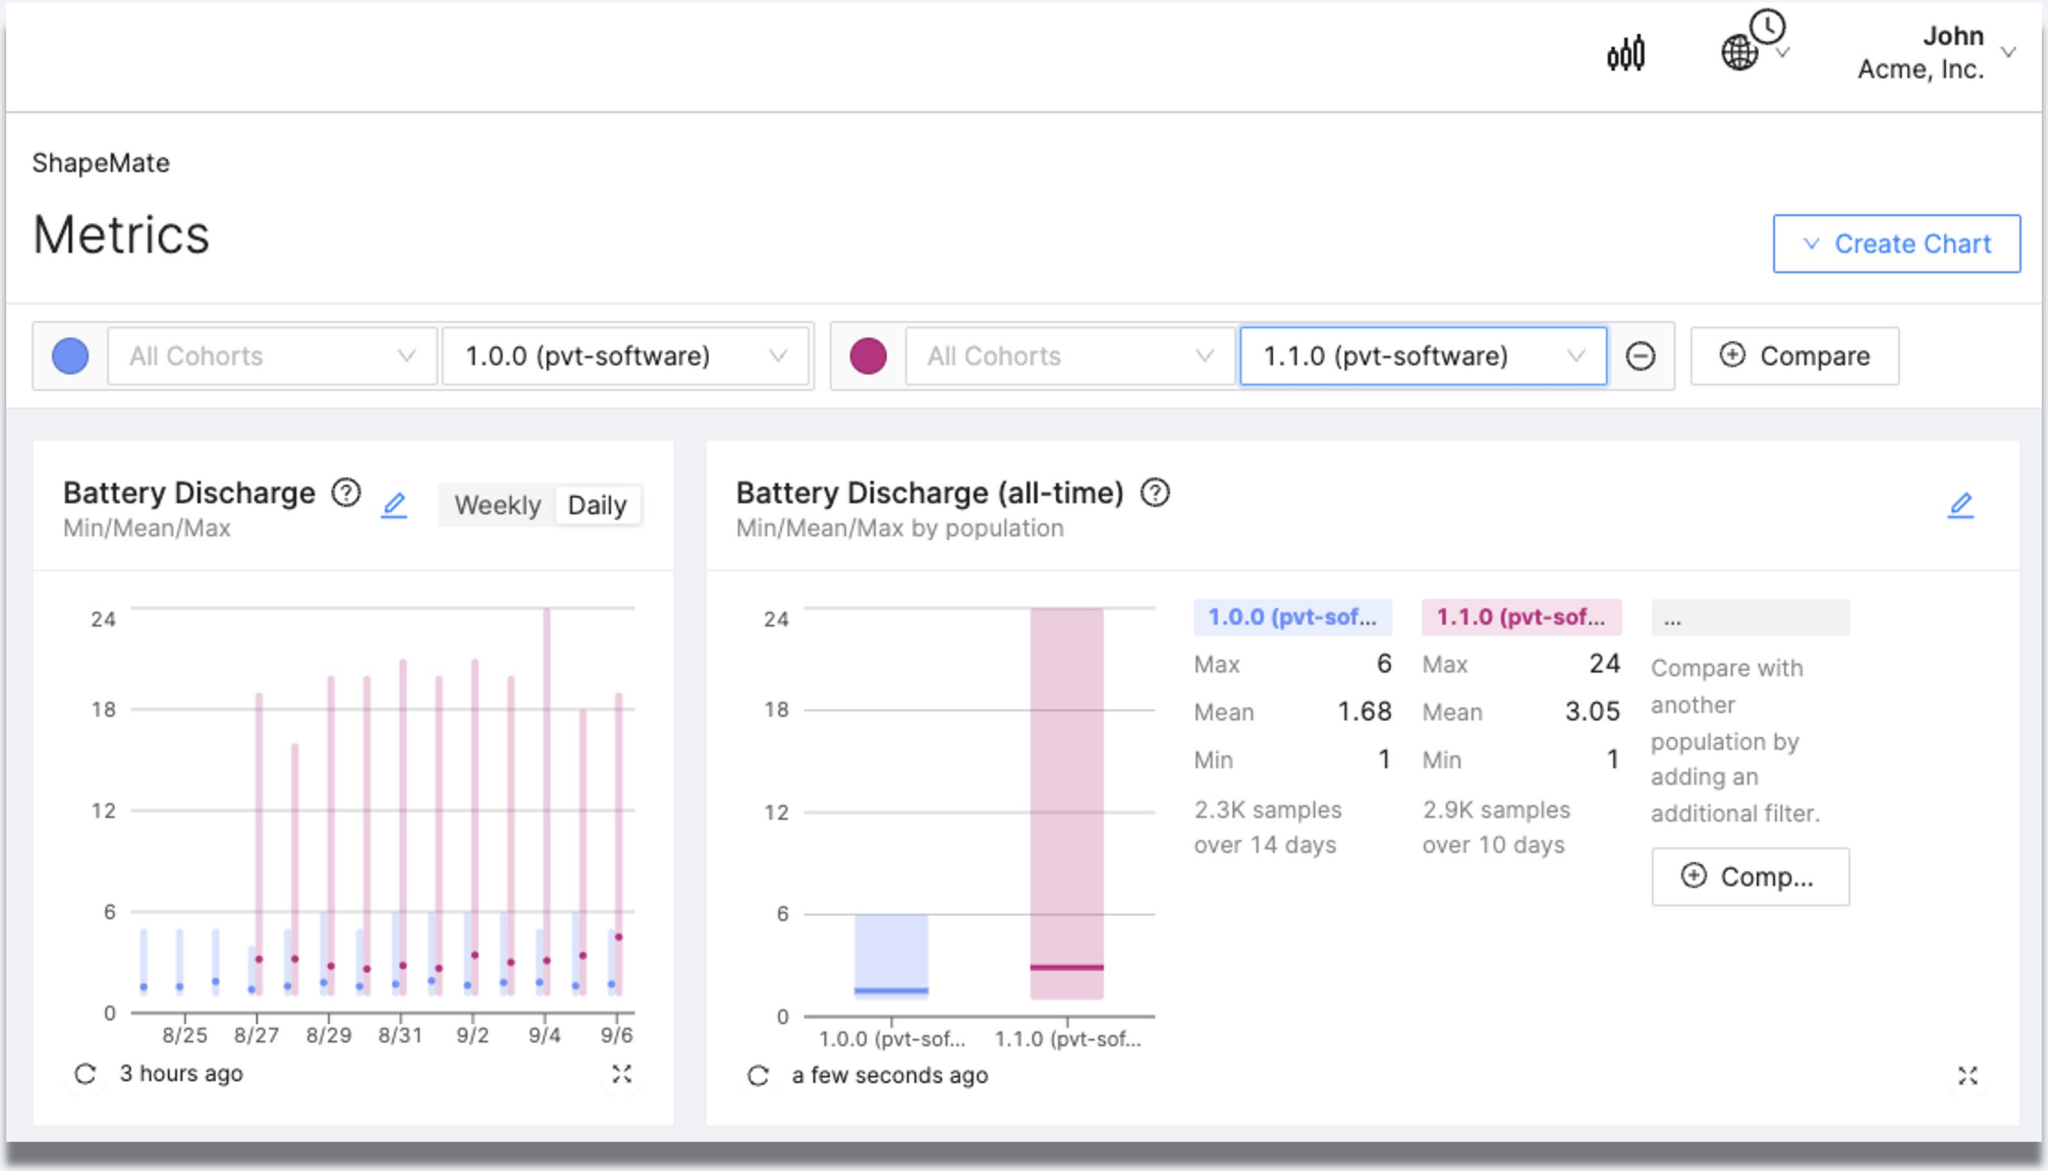

Visualizing Fleet Data

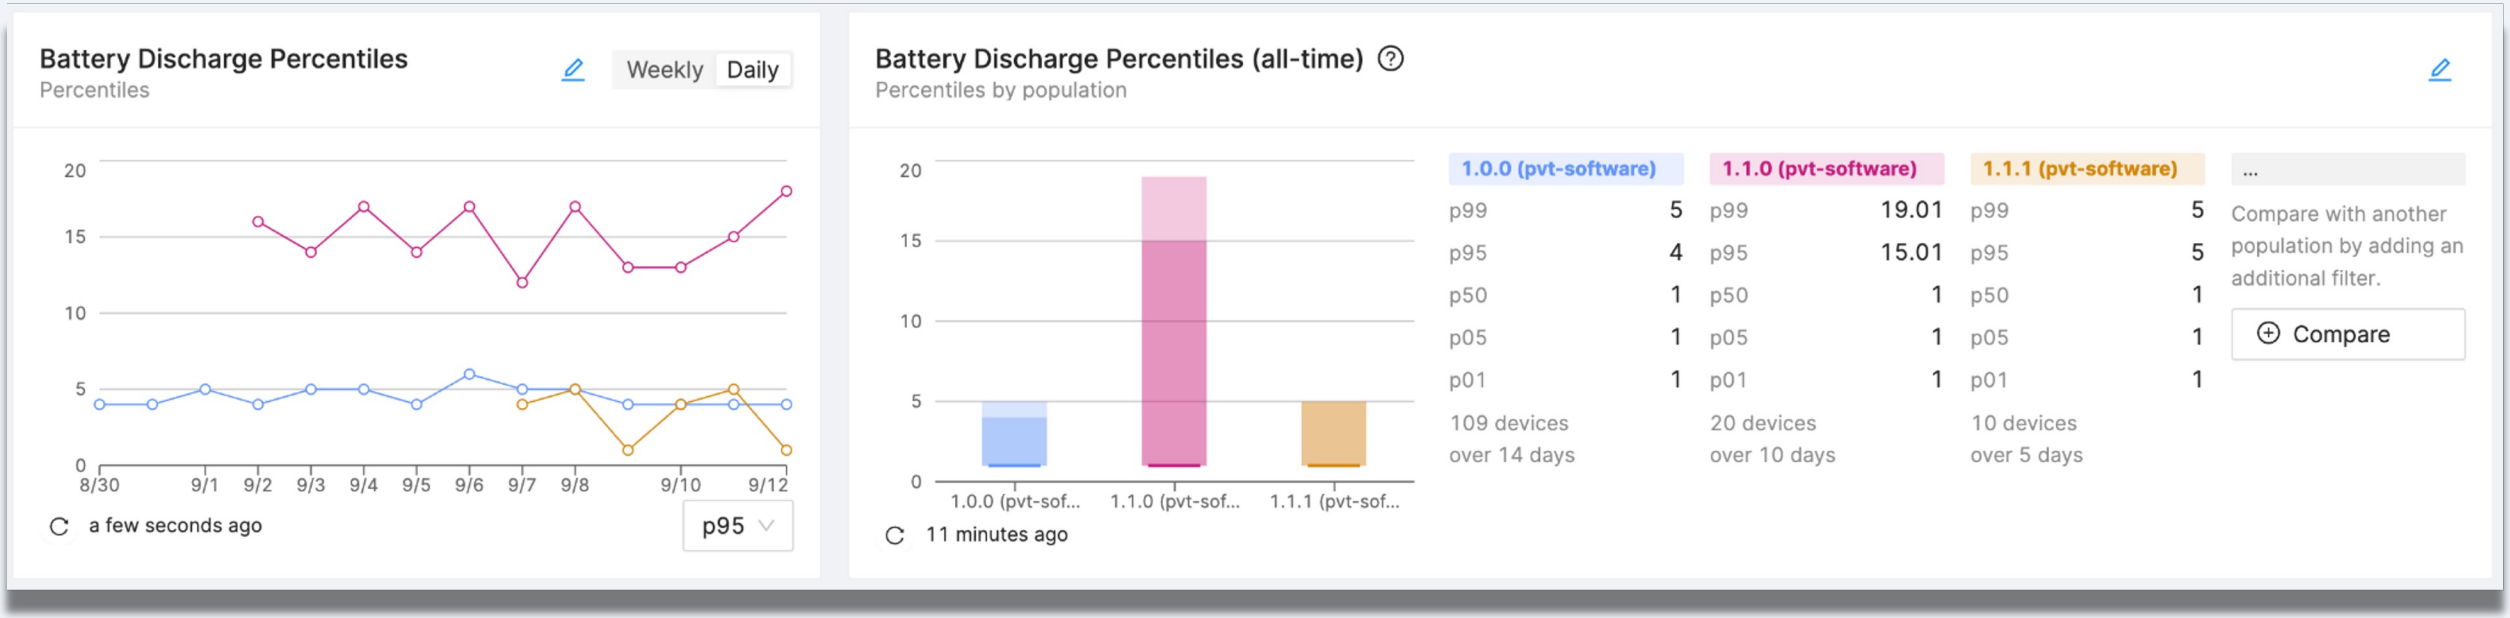

Percentiles help understand how large a problem is and how much of the fleet is impacted.

Normalizing data

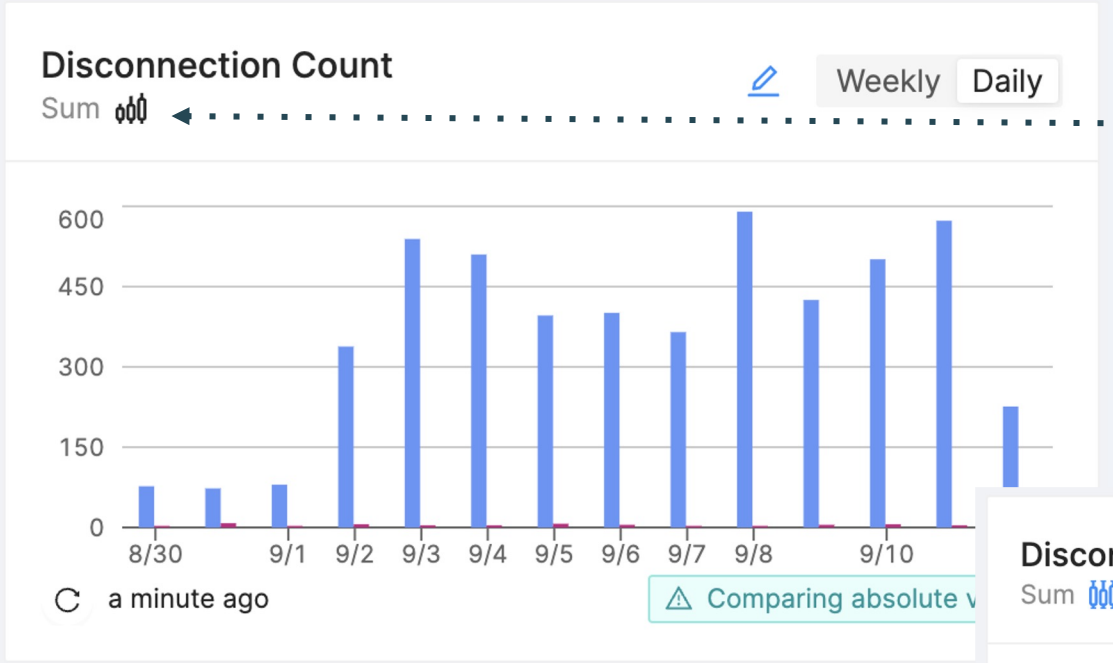

Turn on data normalization to compare absolute values across cohorts of different sizes

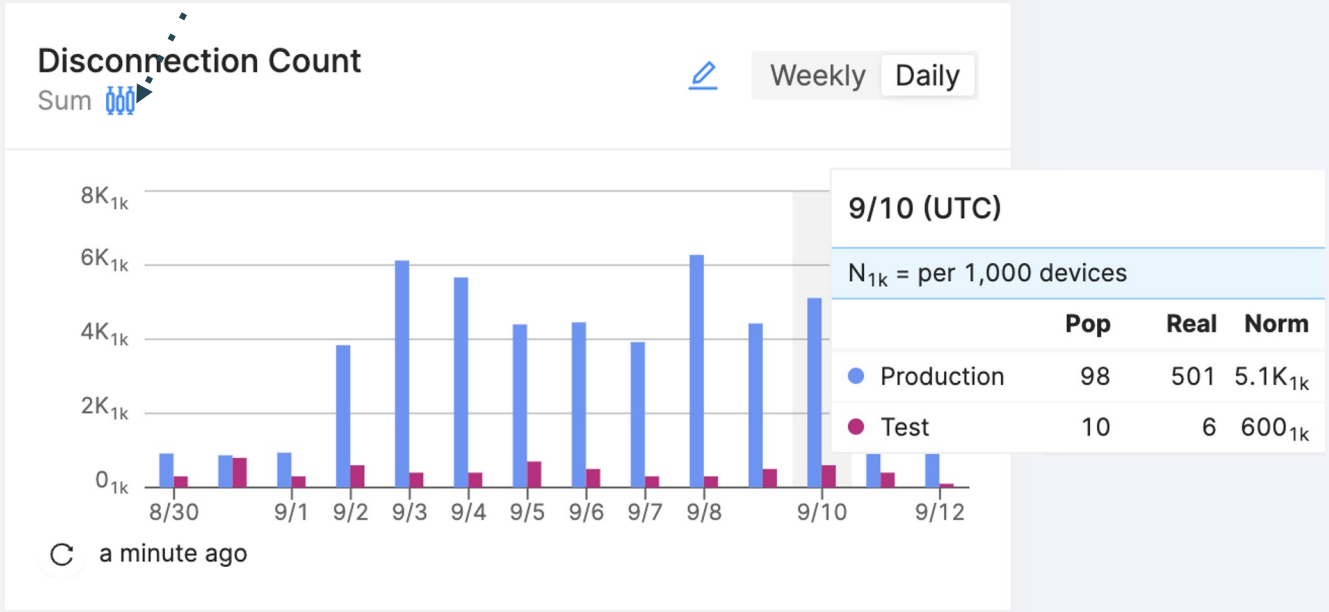

Fleet wide alerts

- Use “Fleet Alerts” to monitor a specific metric over the entire fleet
- Send notification by email and slack to the team

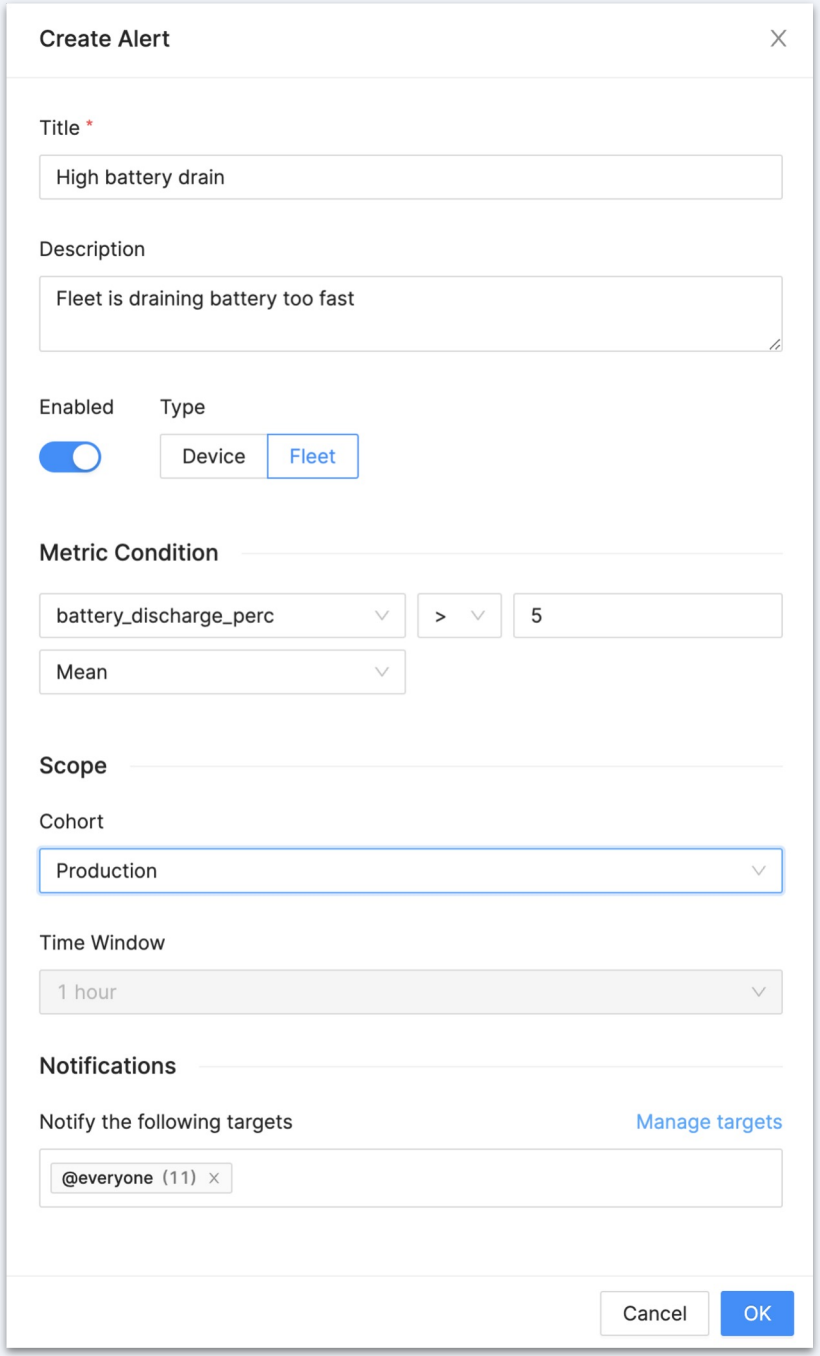

Monitoring Challenges

On Device

- Collecting from different sources and languages
- Partial connectivity
- Flash wear and networking costs

Easy to use “fire and forget” metrics API
On device buffering and aggregation

Backend

- Scaling pains
- Lack of flexibility
- Visualization tools

Usage

- Drowning in data
- Metrics are meaningless when aggregated
- Signal lost in the data

Use best practices to select useful variables and iterate
Use normalization, percentiles, etc

Debugging devices

Poll #2

Where are most of your bugs?

A. Kernel

B. Kernel Modules / Drivers

C. System Daemons

D. Libraries and Runtimes

E. Application Code

Challenges

Poor quality bug reports

Triaging bug reports

No access to the device

No visibility inside the device

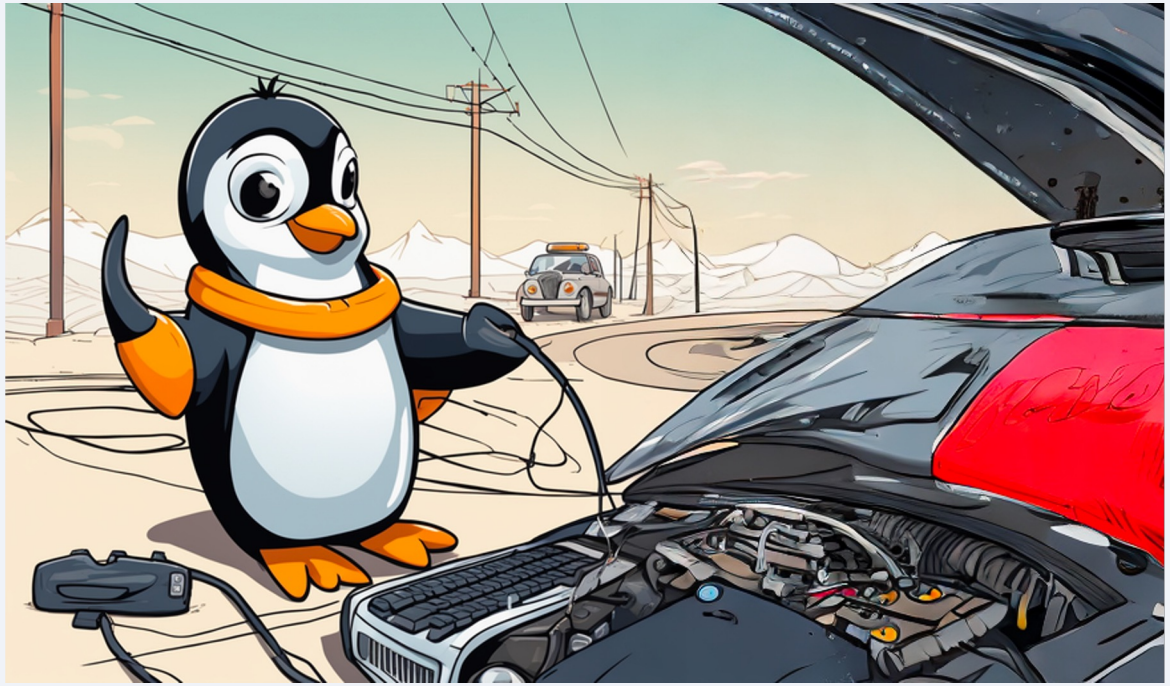

Debugging Tools

- Device Metrics
- Logs
- Coredumps

Logs

Memfault captures device logs with fluent-bit and stores them locally on device.

By default logs are only uploaded for development devices.

You can selectively collect logs from specific devices. This works retroactively.

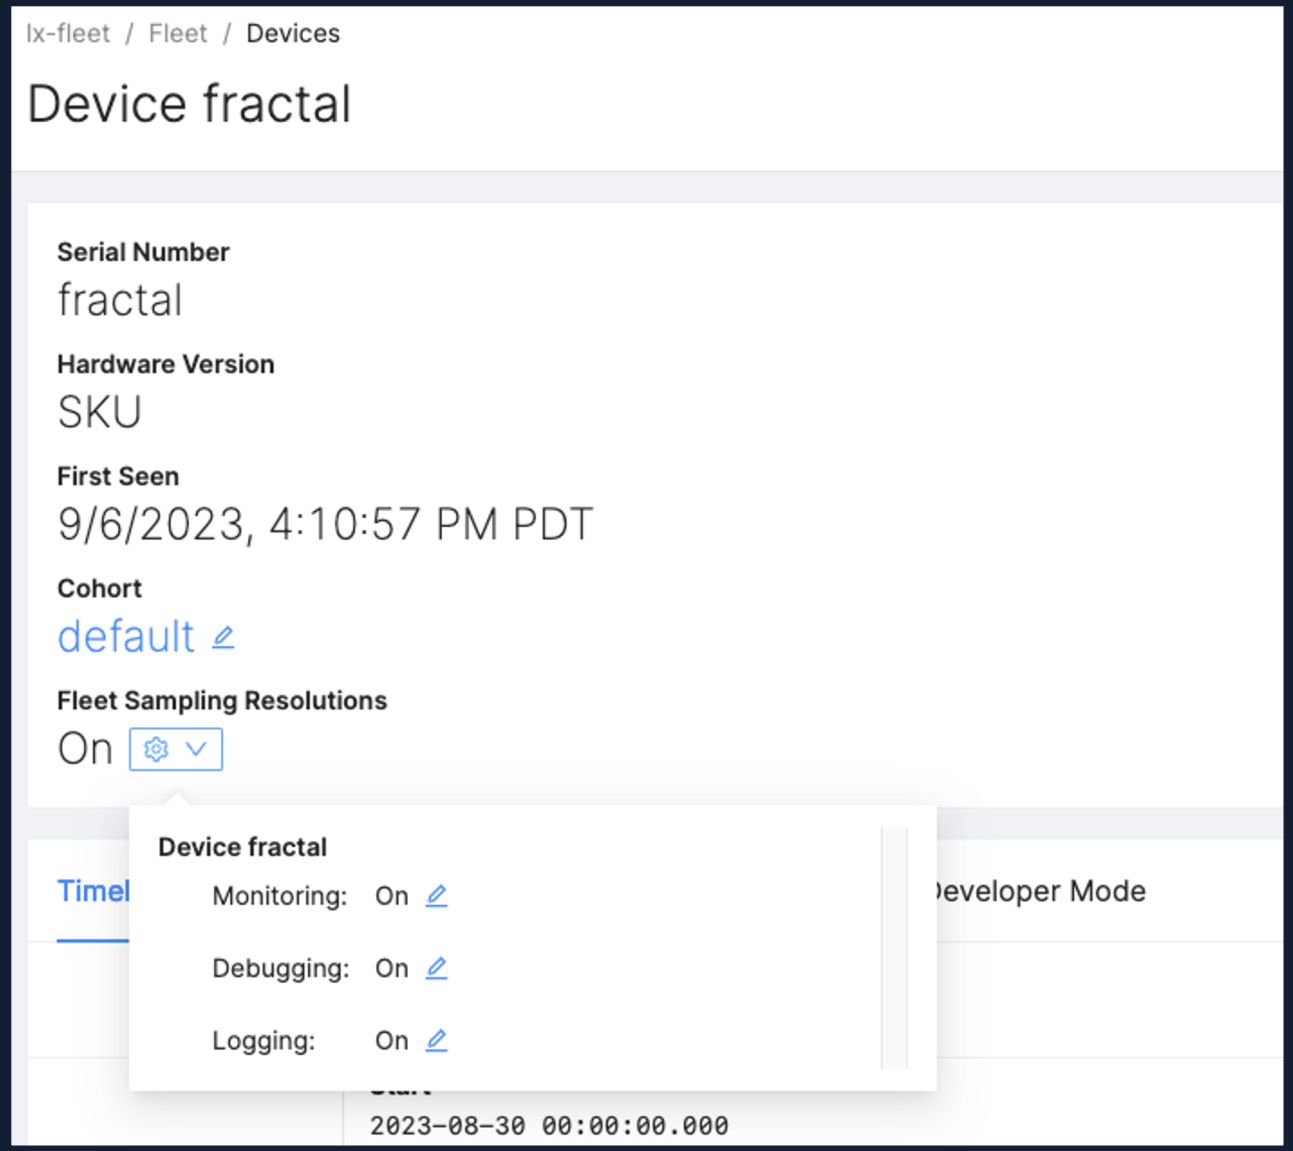

Configuring fluent-bit

Many input options: systemd, files, network, serial

On “the edge” filtering to reduce noise

fluent-bit does not write to disk Sends to memfaultd

```
[INPUT]
        Name systemd

[FILTER]
        Name grep
        Match *
        # Kernel log messages are already forwarded by journald
        Exclude _SYSTEMD_UNIT busybox-klogd.service

[OUTPUT]
        Name tcp
        Host 127.0.0.1
        Port 5171
        Format msgpack
        Match *
        net.keepalive on
        net.keepalive_idle_timeout 10
        # Default retry limit is 1. We recommend setting to a higher value to
        # decrease the chance of losing logs in the event that memfaultd is
        # (re)starting while fluent-bit is attempting to flush logs:
        Retry_Limit 5
```

memfaultd applies rate limiting, compresses and writes to disk at regular intervals (10MB or 1hour)

Logs are kept on disk until requested by backend or max storage space is exhausted.

Browsing logs

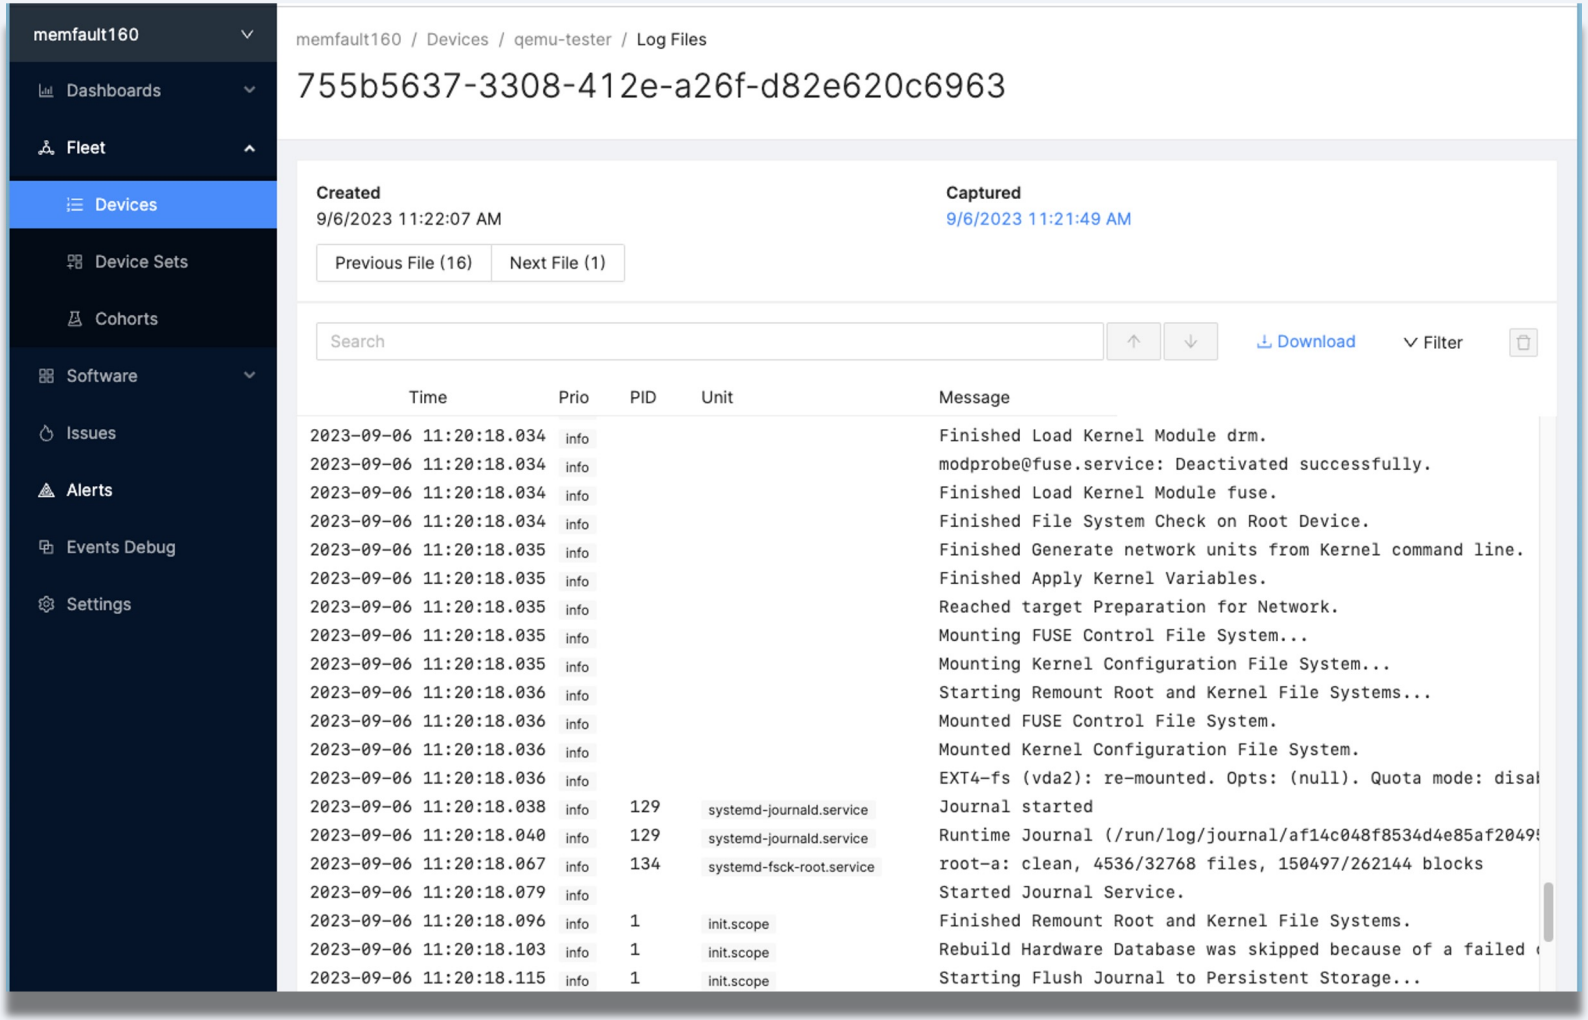

Logs are also available via an API endpoint

Debugging with coredumps

```
 $./test
Segmentation fault (core dumped)
```

Signal received

Action taken by kernel

Coredump

```
CORE(5)                                Linux Programmer's Manual                                CORE(5)

NAME
       core - core dump file

DESCRIPTION
       The  default action of certain signals is to cause a process to terminate and produce a core dump file, a file
       containing an image of the process's memory at the time of termination.  This image can be used in a  debugger
       (e.g.,  gdb(1))  to  inspect  the  state of the program at the time that it terminated.  A list of the signals
       which cause a process to dump core can be found in signal(7).
```

Coredumps contain:

- program status for each thread (registers incl. PC and SP)
- some of the program memory (mostly stack and heap)
- build-id of the running binary
- build-id and address of all the dynamic libraries that are loaded

Using coredumps

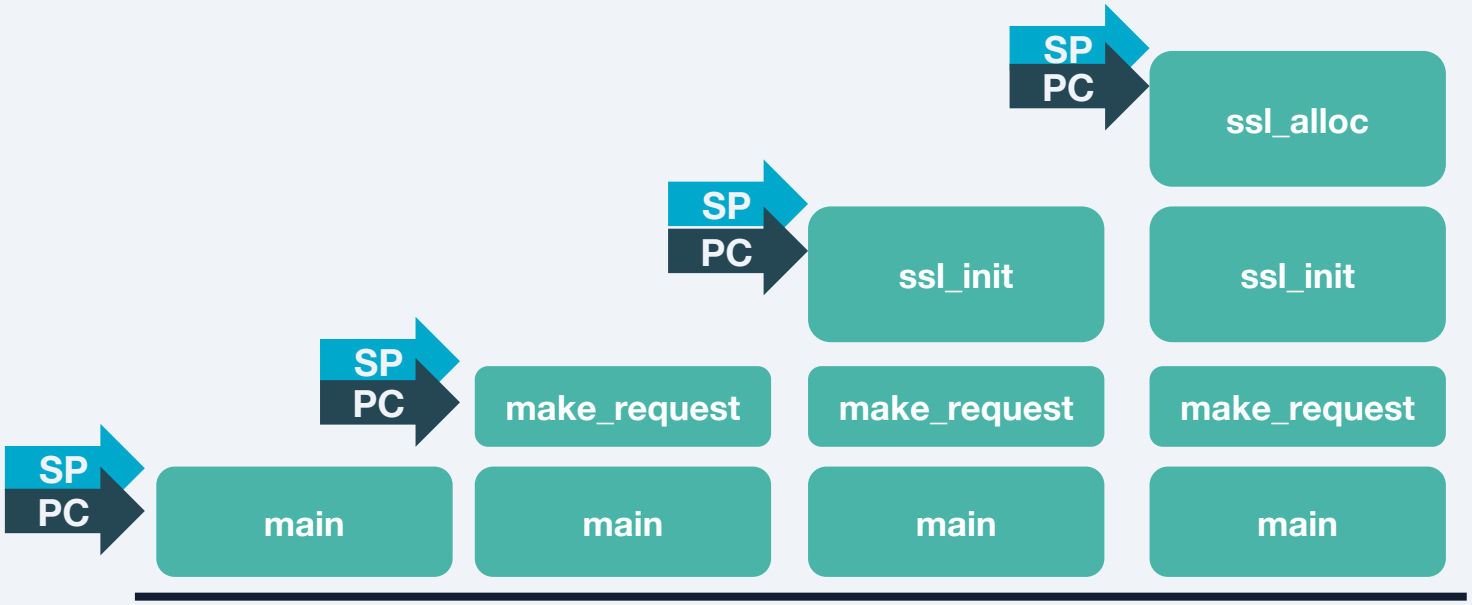

what happens on the device

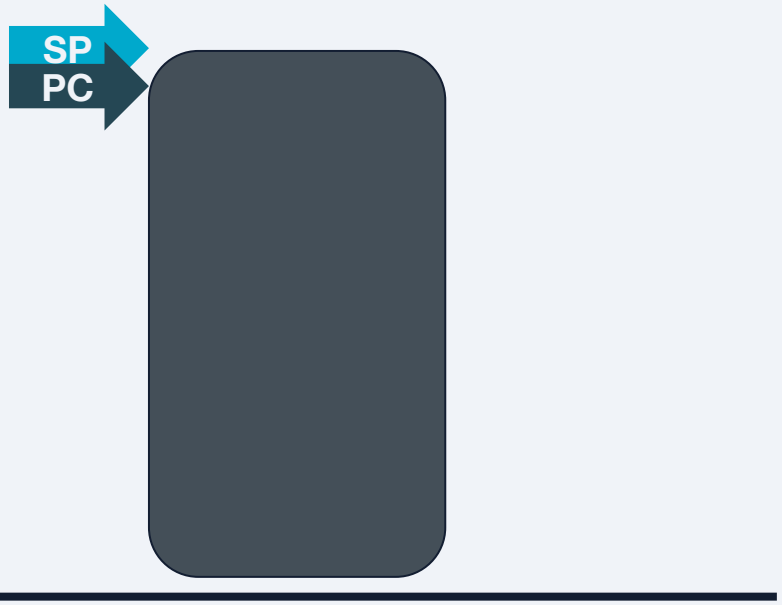

data we get
in the coredump

Using coredumps

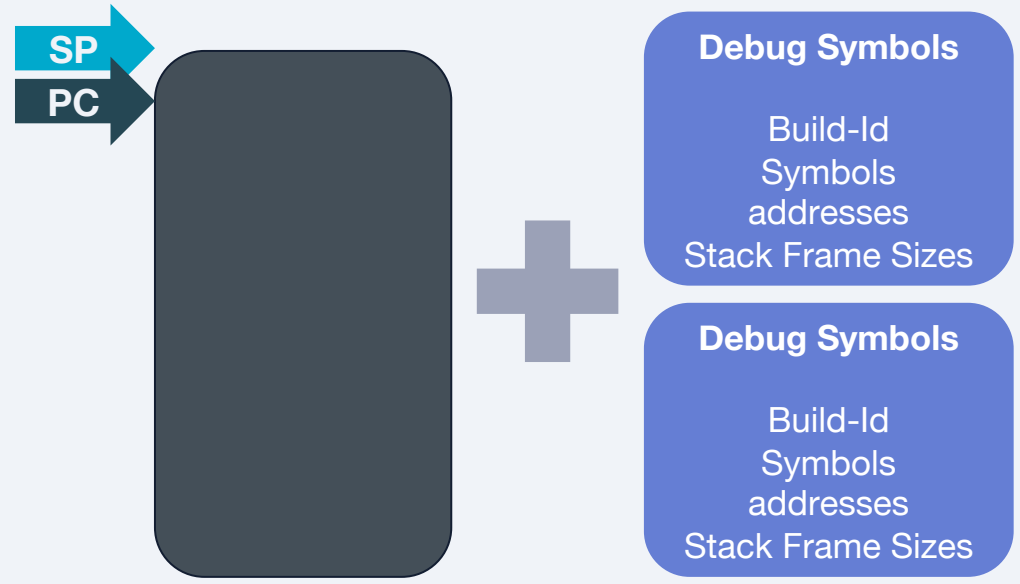

what we get

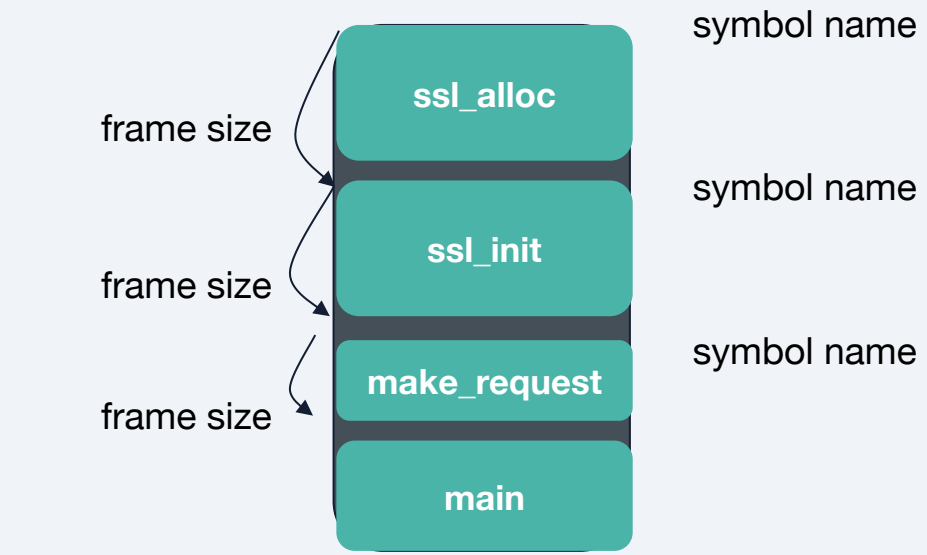

what we can recover

Coredumps view

The state view will list all the running threads and their status

Typically, you will get an immediate read on where the error happened and what were the local variables at the crash

For more complicated bugs, you can download the coredump and run the debugger locally

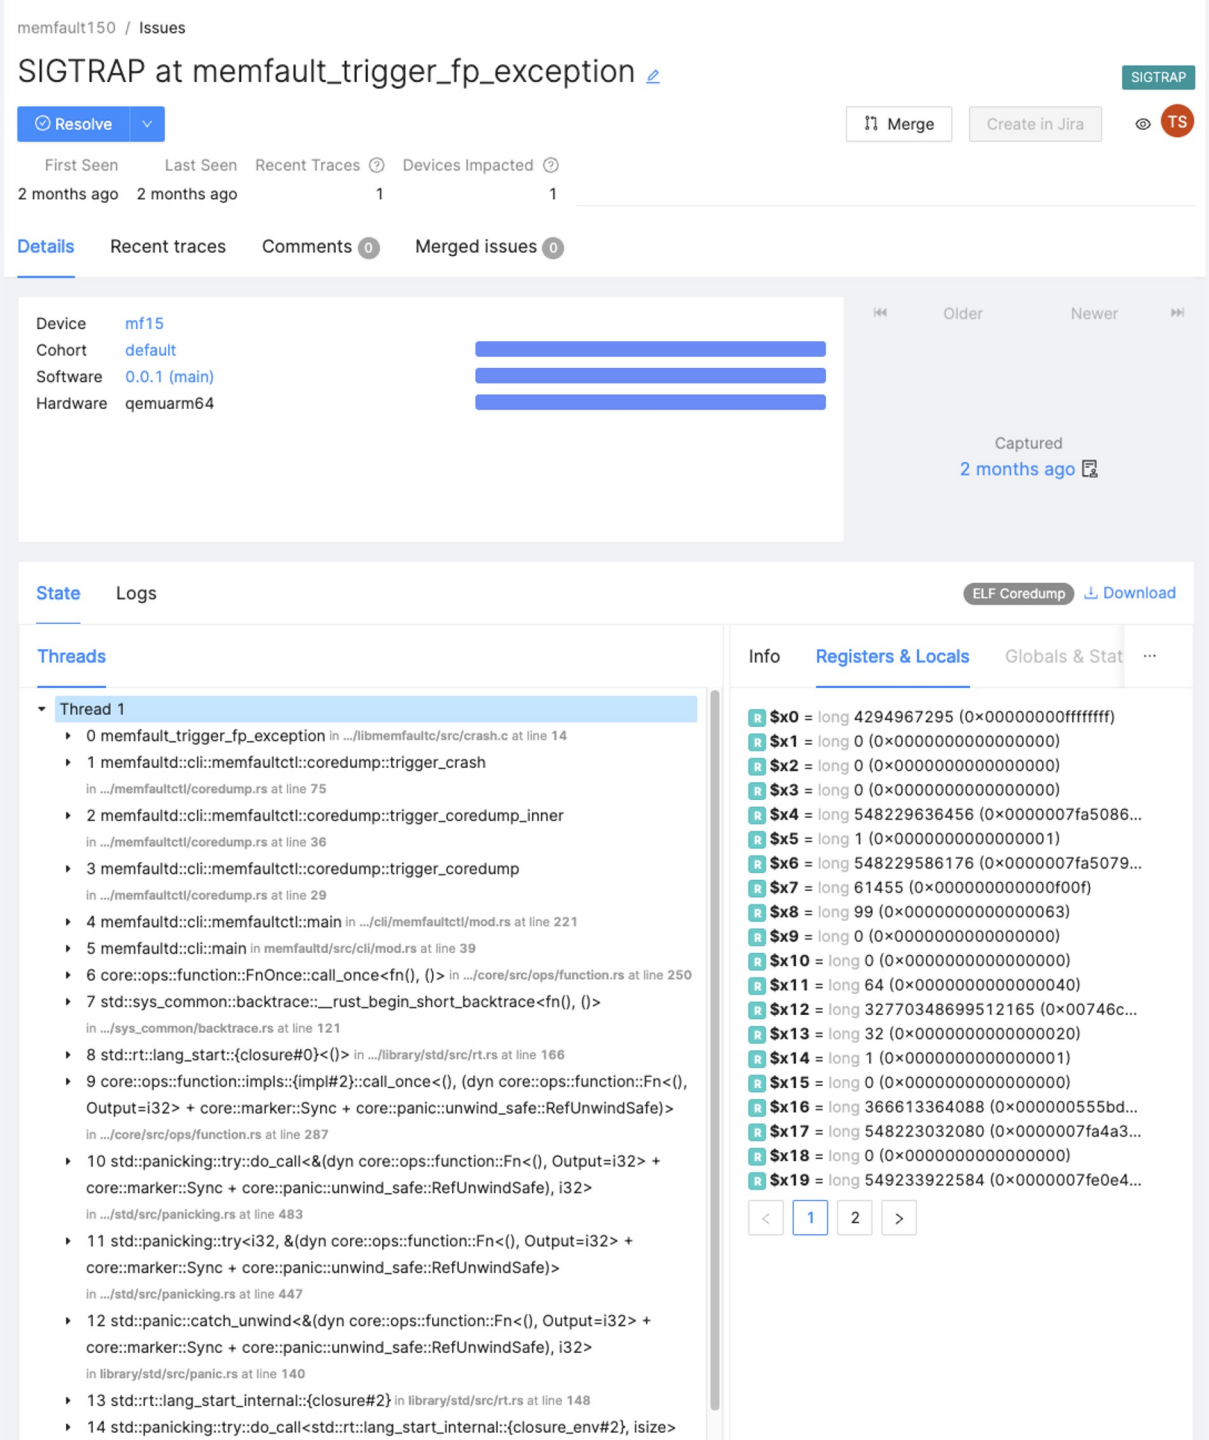

Uploading Symbols

Debugging symbols are required to provide useful coredump analysis

Build and save debugging symbols for all the binaries you produce

Incuding all system libraries

Strip your binaries before sending them to customers

```
# Manual symbols upload

$ gcc -g -o code code.c
$ memfault upload-symbols code
$ strip code
$ cp code /mydevice/usr/bin/


# With Yocto

$ cat >> conf/local.conf
DEPENDS:append = " elfutils-native"
IMAGE_GEN_DEBUGFS = "1"
IMAGE_FSTYPES_DEBUGFS = "tar.bz2"

$ bitbake image
$ memfault upload-yocto-symbols ...
```

Using Coredumps at Scale

Memfault will automatically generate a signature for each **Trace** and group all similar traces together in one **Issue**

Keep track of

- Number of traces captured per day
- Number of devices impacted by a specific issue
- Frequency of an issue over different firmware versions

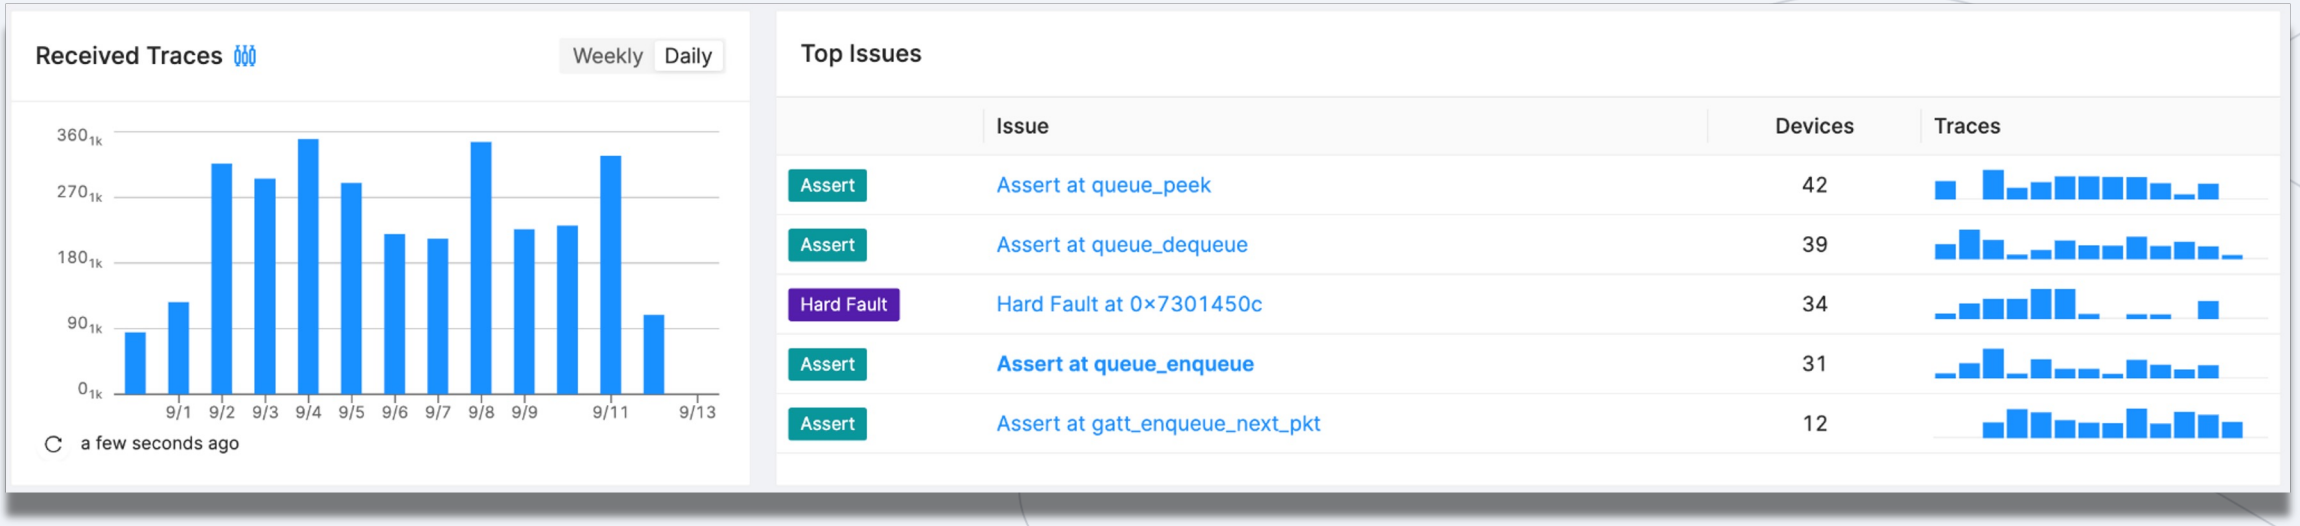

Getting started

Try this at home!

https://docs.memfault.com/docs/linux/quickstart

Memfault Linux SDK

- Docker container to easily build Yocto images
- Pre-configured for OTA with SWUpdate and U-Boot
- Runs inside QEMU or on RasperryPis

```
dev$ git clone git@github.com:memfault/memfault-linux-sdk.git
dev$ cd memfault-linux-sdk/docker
dev$ export MEMFAULT_PROJECT_KEY=abcdef
dev$ ./run.sh -b
docker$ bitbake memfault-image
...
docker$ bitbake swupdate-image
docker$ q


U-Boot 2022.01 (Jan 10 2022 - 18:46:34 +0000)

DRAM:  512 MiB
Flash: 64 MiB
In:    pl011@9000000
Out:   pl011@9000000
Err:   pl011@9000000
Net:   eth0: virtio-net#32
Loading Environment from FAT... OK
Hit any key to stop autoboot:  0
...
```

Thank You!

- memfault.com
- twitter.com/memfault
- interrupt-slack.herokuapp.com

- We’re hiring!

Memfault